

# FASE: Finding Amplitude-modulated Side-channel Emanations

Robert Callan   Alenka Zajić   Milos Prvulovic  
Georgia Institute of Technology

rcallan@gatech.edu, alenka.zajic@ece.gatech.edu, milos@cc.gatech.edu

## Abstract

*While all computation generates electromagnetic (EM) side-channel signals, some of the strongest and farthest-propagating signals are created when an existing strong periodic signal (e.g. a clock signal) becomes stronger or weaker (amplitude-modulated) depending on processor or memory activity. However, modern systems create emanations at thousands of different frequencies, so it is a difficult, error-prone, and time-consuming task to find those few emanations that are AM-modulated by processor/memory activity.*

*This paper presents a methodology for rapidly finding such activity-modulated signals. This method creates recognizable spectral patterns generated by specially designed microbenchmarks and then processes the recorded spectra to identify signals that exhibit amplitude-modulation behavior. We apply this method to several computer systems and find several such modulated signals. To illustrate how our methodology can benefit side-channel security research and practice, we also identify the physical mechanisms behind those signals, and find that the strongest signals are created by voltage regulators, memory refreshes, and DRAM clocks. Our results indicate that each signal may carry unique information about system activity, potentially enhancing an attacker's capability to extract sensitive information. We also confirm that our methodology correctly separates emanated signals that are affected by specific processor or memory activities from those that are not.*

## 1. Introduction

Side-channels are a powerful class of attacks that avoid traditional access controls and protections by exploiting physical or microarchitectural side-effects of computation rather than the computation's overt (algorithm or ISA specified) functionality. Side-channels infer secrets (e.g. cryptographic keys) by observing power consumption [8, 20, 25, 27], sound [5, 12, 33], electromagnetic (EM) emanations [2, 17, 24], behavior under faults [9, 19], and performance of shared caches [6, 36, 38], instruction caches, branch predictors [1], etc.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ISCA'15, June 13-17, 2015, Portland, OR USA

Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3402-0/15/06 \$15.00

DOI: <http://dx.doi.org/10.1145/2749469.2750394>

Side channels that exploit shared hardware resources are well understood in the microarchitecture community, and so proposed hardware and software solutions can alleviate or even completely close such side channels. Side channels based on physical side-effects (power consumption, sound, or EM emanations) are more difficult for microarchitects and programmers to alleviate, in part because the relationship between computational behavior and the resulting side channel signal is very complex and poorly understood. EM emanations side channels may be the most complex: the emanated signals may theoretically be anywhere in the EM spectrum, and signals at different frequencies may provide attackers with insight into different aspects of computational activity.

Unfortunately, the EM side channel is also among the least risky for attackers because EM emanations can be covertly recorded from a distance without modifying, connecting to, or even accessing the victim's system. In contrast, exploits based on shared-resource side channels require attackers to run their "snooper" code in the same system as the victim programs. Power analysis attacks require power measurement equipment to be attached to the system's power supply, preferably at a point close to the processor where the rapid signal fluctuations are not filtered out.

Attack and mitigation efforts must first identify signals that have some dependence on the secret information of interest. For example, differential power analysis attacks [25] analyze the power consumed during encryption operations to find time periods with the greatest dependence on a particular secret key bit. Many EM attacks identify a range of frequencies where EM emanations depend on a secret key bit, then demodulate the signal at those frequencies or filter out unusable frequencies [16, 28, 34]. Many side channel attack descriptions only briefly or implicitly address the underlying mechanisms that cause information leakage because finding information carrying signals and determining their causes are separate processes, and because secret information can be extracted without knowing what causes the information leakage.

Efficient mitigation does, however, require causation. Without a systematic approach to identification and causation, the process of finding root causes is time-consuming and mostly trial-and-error. The defender makes an educated guess about the leakage source, fixes the hypothesized problem, and sees whether the leakage has been reduced.

Systematic instruction-level causation analysis for EM signals caused by direct emanations [3] has recently been devel-

---

This work has been supported, in part, by NSF grant 1318934 and AFOSR grant FA9550-14-1-0223. The views and findings in this paper are those of the authors and do not necessarily reflect the views of NSF or AFOSR.

oped [11]. However, that methodology depends on generating periodic activity at a frequency where precise measurements can be performed, i.e. where there is no interference from external signals such as radio transmitters, or from system-generated noise and interference. Numerous periodic computer system activities are modulated by program activity but have a period that is independent of program activity. In addition, identification of periodic signals that are modulated by specific types of program activity (e.g. LLC misses) has thus far remained an open problem.

This paper describes a methodology we call FASE (Finding Amplitude-modulated Side-channel Emanations). FASE systematically and efficiently identifies periodic signals whose amplitudes change as a result of specific changes in system activity, i.e. signals that are amplitude-modulated by system activity. Our methodology uses micro-benchmarks to generate repetitive changes in processor and memory activity, then processes the resulting EM signals to find spectral patterns corresponding to amplitude modulation. The EM spectrum is full of amplitude-modulated signals (e.g. radio broadcasts) that are not modulated by program activity. FASE filters out such signals by generating several different activity patterns and reporting only those signals which are specifically modulated by all the generated activity patterns.

Our FASE approach for discovering AM-modulated signals is highly effective. Our experiments cover the entire AM radio spectrum, and were performed without shielding in a major metropolitan area with hundreds of radio stations nearby. Furthermore, computer systems produce thousands of periodic signals that are not modulated by system activity. FASE successfully rejected all such signals, while reporting the small number of remaining signals that were indeed modulated by the tested system activities. We validated the automated FASE procedure through manual inspection of all rejected signals that were similarly strong (or stronger) than the FASE-reported ones, confirming that these rejected signals do not measurably respond to changes in system activity.

To demonstrate the usefulness of FASE and to understand potential EM side-channel vulnerabilities of modern processor and memory systems, we then identified the source of each periodic signal and the mechanism by which it was modulated. We first identified the source of each signal using short-range probes to find the apparent origin of the signal within the system. Then we examined data sheets of the components nearby to explain how the periodic signal is generated. Finally we performed additional micro-benchmark experiments to identify the modulation source.

We discovered three main types of signals. First, strong signals emanate from switching voltage regulators and power filtering components at the specified switching frequency of the regulator (usually between 200kHz and 500kHz) or multiples of it (harmonics). These signals are modulated by variations in power consumption in the voltage regulator's load (the processor, memory or other system components), and

they allow attackers to carry out the equivalent of power side-channel attacks from a distance without the need to place probes within the system. Voltage regulators for processor cores, the memory controller, and the DRAM memory itself often have different switching frequencies, giving the attacker component-by-component power consumption information.

Another type of signal is generated by periodic memory refreshes. This signal is amplitude-modulated by memory access activity, i.e. the attacker gets an at-a-distance readout of how often the memory is used. Unlike voltage regulators, which can be considered an external problem by processor/memory architects, these refresh-related signals are entirely caused by activity within the purview of memory controller designers and are likely to be completely eliminated by appropriate modifications to how memory refresh is carried out.

At higher frequencies, FASE discovers clock signals and their harmonics that are modulated by activity in the clock's domain. Because most clock and switching regulator harmonic frequencies are subject to electromagnetic interference (EMI) regulations [15], they are subjected to measures (such as spread-spectrum clocking) that spread the resulting EM emanations over a range of frequencies [22]. In spite of this, FASE discovers such signals and provides insight into the nature of the activity that modulates them. In particular, we identify that DRAM clocks generate EM emanations which are modulated by DRAM activity. The systems tested generated weak spread-spectrum signals at CPU clock frequencies. Interestingly, we do not observe any variation in these signals in response to processor activity.


FASE can be used to find which parts of a system leak information about some aspect of program activity. This can be used to reduce the strength of modulated signals and to weaken their modulation, i.e. disrupt the connection between program behavior and the variations in activity that modulate such signals. Using memory refresh signals as an example, this would involve randomization of the interval between refresh commands, while modulation-weakening efforts might involve careful scheduling of memory accesses to avoid their interaction with refresh activity.

## 2. Methodology for FASE

### 2.1. Overview of AM-Modulation in Computer Systems


The spectral properties of amplitude modulated signals are well-understood [32], but unintentional AM signals in computer systems have some properties not typically found in traditional uses of AM signals (i.e. telecommunications). To understand why FASE is needed and how it uses generated modulation patterns to identify AM-modulated signals, a review of the general properties of AM modulation and the irregularities of "accidental" side-channel transmission is needed.

Figure 1 shows the spectrum of an ideal carrier signal (at frequency  $f_c$ ) that is modulated by an ideal sinusoidal signal at frequency  $f_{alt}$ . In addition to the carrier signal, this spectrum has strong "side-band" signals offset by  $f_{alt}$ , i.e. at frequencies




**Figure 1: Sinusoidal carrier modulated by a sinusoidal signal.**

$f_c - f_{alt}$  and  $f_c + f_{alt}$ . This would be the spectral pattern to look for when a periodic signal has a perfectly stable frequency and is modulated by a pattern of activity with a fixed period of  $T_{alt} = 1/f_{alt}$  with no variation in timing but these ideal conditions are rarely present in unintentional signals.



**Figure 2: Sinusoidal carrier modulated by an arbitrary signal.**


Figure 2 shows the spectrum of an ideal sinusoidal carrier modulated by a realistic signal, such as program activity. The two side-band signals now correspond to the spectrum of the modulating activity. The tallest spike in each side-band signal corresponds to the dominant periodic behavior of that activity and the smaller “bumps” in each side-band signal indicate other common periods of repetitive activity. This type of non-ideality is typical for program-generated repetitive behavior: for a given task, the time each repetition of the task takes is not always the same, but there are often several commonly-occurring execution times among the repetitions. For example, in multi-processor or SMT systems the repetitions of a loop may take longer or shorter depending on timing variations due to resource contention with other running threads.



**Figure 3: Non-ideal carrier modulated by a sinusoidal signal.**


Figure 3 shows a non-ideal carrier modulated by an ideal signal. The spectrum for the carrier is now spread around its nominal value and this spreading is also present in the two side-band signals. Even though  $f_{alt}$  is perfectly stable, the side-bands at  $f_c - f_{alt}$  and  $f_c + f_{alt}$  will “inherit” the instability of  $f_c$ . Many periodic signals are spread out in this manner in

computer systems. For example, spread-spectrum clocking results in deliberate spreading of the clock signal’s frequency. Additionally, many periodic activities (e.g. voltage regulator switching) do not require precise timing, so they often use less stable (cheaper/simpler) oscillators. Combining a non-ideal carrier (Figure 3) with a non-ideal modulating activity (Figure 2) produces the spectrum in Figure 4.



**Figure 4: Non-ideal carrier modulated by an arbitrary signal.**

Several other non-ideal properties of computer systems are manifested in measured spectra. Randomly timed switching activity causes broadband noise, and this noise appears as gently rolling “hills” and “valleys” in the spectrum. Additionally, a realistic spectrum contains periodic signals from both inside and outside the system that are either not modulated at all or that *are* AM-modulated (e.g. AM radio broadcasting) but not by program activity. Such a spectrum is shown in Figure 5. Even if we know the carrier and the program activity’s frequency content it is hard to decide whether this spectrum contains an activity-modulated signal by visual inspection. Our FASE methodology uses several specially generated program activities in conjunction with a heuristic carrier likelihood function to automate the decision process and overcome these problems.



**Figure 5: The same non-ideal carrier and arbitrary side-band signal as Figure 4 with noise and other sources present.**

Many periodic carrier signals in computer systems are generated by digital circuits and clocks, and therefore have sharp transitions that are best approximated by rectangular pulses instead of the sinusoidal waves used as carriers in communications systems. The spectrum of a pulse train with an arbitrary duty cycle is equivalent via Fourier analysis to a set of sinusoids with various amplitudes at  $f_c$  and its multiples (harmonics). In other words, for each carrier signal generated by a digital circuit or clock, additional carrier signals will also be present at  $2f_c$ ,  $3f_c$ ,  $4f_c$ ,  $5f_c$ , etc. As the duty cycle of a signal approaches 50%, the amplitudes of the odd-numbered

harmonics ( $f_c$ ,  $3f_c$ ,  $5f_c$ , etc.) reach their maximum, while amplitudes of the even harmonics ( $2f_c$ ,  $4f_c$ , etc) trend toward zero. For a small duty cycle (i.e.  $< 10\%$ ), the magnitudes of the first few harmonics (both even and odd) decay approximately linearly. Finally, note that these observations imply that the amplitudes of all the harmonics are a function of the duty cycle. If program activity modulates the duty cycle of a periodic signal while keeping its period constant (i.e. causes pulse width modulation), all of the signal's harmonics are amplitude-modulated and consequently will be identified by our FASE methodology.

## 2.2. Creating System Activity at Controlled Frequencies

A carrier at frequency  $f_c$  modulated by system activity is a lot easier to recognize if we generate periodic processor and/or memory activity that repeats  $f_{alt}$  times per second. Micro-benchmarks for generating such periodic activity have already been proposed for demonstrating the presence of EM emanations from computer systems [39] and for creating measurable periodic signals at arbitrary frequencies [11]. A simplified version of one such micro-benchmark is shown in Figure 6. The loop beginning on line 2 performs one activity (activity X), and the loop beginning on line 8 performs another activity (activity Y). The outer loop repeatedly alternates activities X and Y, creating periodically changing activity whose period equals the execution time for one iteration of the outer loop. This alternation period  $T_{alt}$  is the inverse of the frequency  $f_{alt} = \frac{1}{T_{alt}}$ . Note that prior uses of similar micro-benchmarks [11, 39] used this alternation to *generate a carrier signal* at some chosen frequency  $f_c$ , while we use this alternation at  $f_{alt}$  to measure AM-modulation of any potential *carrier signals intrinsically generated* (and emanated) by the system.

```

1  while(true){
2    // Execute the X activity
3    for(i=0;i<inst_x_count;i++){
4      ptr1=(ptr1&~mask1)|((ptr1+offset)&mask1);
5      // The X-instruction, e.g. a load from L2
6      value=*ptr1;
7    }
8    // Execute the Y activity
9    for(i=0;i<inst_y_count;i++){
10     ptr2=(ptr2&~mask2)|((ptr2+offset)&mask2);
11     // The Y-instruction, e.g. a store from L2
12     *ptr2=value;
13   }
14 }
```

**Figure 6: Pseudo-code to generate the X/Y alternation activity.**

As an example of how the alternation of activity can AM-modulate a carrier signal, consider the DRAM clock. Activity X may involve many LLC misses, so it results in substantial DRAM activity. During the X-activity half-period, the DRAM clock drives a lot of switching activity (current flowing through wires), resulting in strong emanations at the DRAM clock frequency. If activity Y has little DRAM activity, less switching

activity is driven by the DRAM clock, generating weaker emanations at the DRAM clock frequency. Therefore the amplitude of the emanations at the DRAM clock frequency will change with period  $T_{alt}$  (frequency  $f_{alt}$ ), which means that emanations at the DRAM clock frequency will be AM-modulated by the X/Y periodic behavior whose frequency is  $f_{alt}$ .


It is important to note that switching between activity X and activity Y is abrupt and that we adjust the `inst_x_count` and `inst_y_count` variables so that activity X and activity Y are each done for half of the alternation period (50% duty cycle). Thus the spectrum of each side-band around the carrier's frequency  $f_c$  will also have odd-numbered harmonics of the alternation frequency, i.e. the side-band signal will have spikes/peaks at  $f_c \pm 3f_{alt}$ ,  $f_c \pm 5f_{alt}$ , etc. in addition to  $f_c \pm f_{alt}$ . Also note that the alternation frequency  $f_{alt}$  can be controlled by changing the loop counts, allowing us to create several spectra (with different  $f_{alt}$ ). These spectra can be considered jointly in an effort to distinguish which carriers are modulated by a particular activity.

FASE results for different X/Y pairings usually provide a strong indication of which aspect of the system modulates a given carrier signal. For example, when a signal at a particular frequency  $f_c$  is modulated by X/Y alternation between memory activity and any on-chip activity, but remains unmodulated when alternating between two types of on-chip activity, the carrier signal and/or its modulation mechanism are likely related to the memory controller, processor-memory communication, or the DRAM memory itself.

## 2.3. FASE: Finding Amplitude-modulated Side-channel Emanations

As indicated in Section 2.1, discovery of activity-modulated carriers by “eyeballing” the spectrum without generating controlled system activity would be very difficult. Theoretically, one could look for narrow spikes (potential carriers) with symmetric side-bands on either side as shown in Figure 2, but this approach is not practical due to the non-ideal nature of unintentional carriers, the interference of other signals, and noise as shown in Figure 5.

Measuring arbitrary programs or benchmarks may provide some information about carriers that are modulated by system activity but it would be difficult to determine the spectral properties of such arbitrary system activity. Even if we are somehow given spectral information about activity in an application, it would be hard to recognize whether the side-band signals around each potential carrier match that spectrum with high confidence because 1) amplitude modulation combines (convolves) the spectrum of the possibly non-ideal carrier signal with the arbitrary benchmark spectrum (Figure 4), and 2) recognition of such a complicated overall spectrum is further hampered by noise and unrelated signals that overlap with portions of the modulated-signal spectrum (Figure 5). We cannot directly control the shape of a system's intrinsic carrier signals, but we can use the micro-benchmarks described in Section 2.2



**Figure 7: A carrier at  $f_c$  and its right and left side-bands generated by memory activity.**


to generate system activity that is as close to a perfect square wave as possible. This results in side-band signals whose spectrum has a shape that closely matches the shape of the carrier signal they are modulating, with a  $f_{alt}$  separation between the carrier and its two side-bands in the spectrum. This could be used to find carriers automatically by looking for such right and left side-band signals because they always appear as peaks in the spectrum separated by  $2f_{alt}$  with the carrier peak half-way between them. However, this simplistic approach has a number of drawbacks. First, the alternation activity is a square wave which has many odd-numbered harmonics ( $f_c \pm f_{alt}$ ,  $f_c \pm 3f_{alt}$ , etc.) that are separated by exactly  $2f_{alt}$ . This makes it difficult to attribute the spikes in the side-band signals to particular carrier frequencies, creating many false positive indications of carrier locations. Second, for some values of  $f_{alt}$ , some of the side-band signals may be overwhelmed by noise and unrelated signals, which would result in many false negatives. Third, computer systems contain many components with periodic activity, so unmodulated signals are often concentrated at specific frequencies. Some such spectral peaks will be nearly  $2f_{alt}$  apart by random chance, resulting in more false positives.

Many of the problems caused by the harmonics of the alternation signal and by the existence of unrelated signals can be solved by performing multiple measurements with different alternation frequencies, e.g.  $f_{alt1}$ ,  $f_{alt2} = f_{alt1} + f_{\Delta}$ ,  $f_{alt3} = f_{alt1} + 2f_{\Delta}$ , etc., where  $f_{\Delta}$  is typically small compared to  $f_{alt}$ . Figure 7 shows five recorded spectra with  $f_{alt1} = 43.3\text{kHz}$  and  $f_{\Delta} = 0.5\text{kHz}$  around a carrier signal at  $f_c = 1.0235\text{MHz}$ . To avoid clutter, Figure 7 only shows the three parts of the spectrum that contain the left side-bands, the carrier, and the right side-bands of the signals. In other words, it does not show about 40kHz worth of spectrum to the left and right of the carrier. Note how the peaks in the side-bands move by  $f_{\Delta}$  as the alternation frequency  $f_{alt}$  changes by  $f_{\Delta}$ .

Conceptually, the FASE methodology for finding activity-modulated carriers and determining the frequencies of such

carriers is now as follows. First, perform several measurements (we use five) with different  $f_{alt}$  frequencies as described above. Second, look for a shape in the spectrum that moves by  $f_{\Delta}$  or  $-f_{\Delta}$  in successive measurements. This approach eliminates external signals and system-emanated periodic signals that do not correspond to activity-induced AM modulation because such signals stay at the same frequency as  $f_{alt}$  changes. It also only detects the first harmonic of  $f_{alt}$  to the right and left of the carrier. Recall that the alternation activity changes abruptly and may not have a perfect 50% duty cycle, so the spectrum of the modulated signal has side-band signals not only at  $f_c \pm f_{alt}$  but also at  $f_c \pm 2f_{alt}$ ,  $f_c \pm 3f_{alt}$ , etc. However, only the first harmonic ( $f_c \pm f_{alt}$ ) moves by  $f_{\Delta}$  in the spectrum as we change  $f_{alt}$  by  $f_{\Delta}$ . The other harmonics in the side-band move by  $2f_{\Delta}$ ,  $3f_{\Delta}$ , etc.

Once we have identified a first harmonic side-band signal in this way, we can determine whether it is the left side-band (moves by  $-f_{\Delta}$ ) or the right one (moves by  $f_{\Delta}$ ), and we can compute the frequency of its carrier signal. The carrier is located at  $f - f_{alt_i}$  if the modulated peaks are detected at frequency  $f$  and if  $f_{alt_1}$  is to the left of  $f_{alt_5}$  (or at  $f + f_{alt_i}$  if  $f_{alt_1}$  is to the right of  $f_{alt_5}$ ). Note that detection of a single harmonic of  $f_{alt}$  in a single side-band is sufficient to detect a carrier frequency, i.e. we do not need all of them to find the frequency of the carrier. Also, note that any harmonic (e.g.  $\pm 2\text{nd}$ ,  $\pm 3\text{rd}$ , etc.) is sufficient since the observed spacing between the side-band peaks is unique for each harmonic (e.g.  $2h_{\Delta}$  for the positive 2nd harmonic,  $-3h_{\Delta}$  for the negative third harmonic, etc.). This comes in handy if one or more of the signals overlap with other signals or unusually strong noise – with five measurements we get a total of ten side-band signals (two side-bands per measurement) at different frequencies, so we can reliably detect the presence of modulation and the frequency of the carrier even if several of the side-band signals are obscured as shown in the left side-band of Figure 12. Also note that this approach does not rely on actually observing a



**Figure 8: Simplified spectrum representation of the harmonics of the LDL2/LDL1 activity for the Intel Core i7 desktop.**

peak for the carrier signal. This is important when the carrier itself is located in a crowded part of the spectrum – as long as at least a few side-band signals “land” in a “quiet” part of the spectrum, we can deduce the exact frequency of the carrier.

There often are several modulated carrier signals in the same general region of the spectrum, so that their side-band signals may not be neatly separated from each other. A simplified representation of one actual recorded spectrum is shown in Figure 8. The thick lines in this figure indicate carrier frequencies, each with a different color. The thin lines indicate the frequencies of side-band  $f_{alt}$  harmonic signals, where the color indicates which carrier generates this side-band signal and the number indicates which harmonic of  $f_{alt}$  it corresponds to. Without FASE the interleaved side-band signals generated by different carriers make it very difficult to manually interpret such measured spectra.

The antennae we used to capture signals from computer systems were designed to detect broadcast radio signals over a wide frequency range, so they pick up these interfering signals very well. It is critical to note that FASE is intended to identify only AM signals which are modulated by *our micro-benchmark*. Although AM radio signals are amplitude-modulated and strong, FASE correctly identifies that these signals are *not caused by our modulation activity* and so should not be reported. This is important not only because it is painfully expensive to shield a measurement setup from broadcast signals, but also because computer systems themselves emit strong radio signals (wifi, bluetooth, NFC, etc.) that are modulated for communication purposes but should not be reported by FASE unless they are *also* modulated by our microbenchmark activity.

#### 2.4. Automating FASE

In Section 2.3, we explained that carriers are found by searching for a shape that shifts by  $f_{\Delta}$  when  $f_{alt}$  changes by  $f_{\Delta}$ . However, visual comparison of numerous recorded spectra across a wide range of frequencies would be tedious and error

prone. Unfortunately, the scope and length of this paper do not allow a detailed description of our side-band detection and analysis algorithms. Instead, we provide insight by presenting a simplified and easily-implementable heuristic for finding side-bands whose shifts in frequency correspond to shifts in  $f_{alt}$ . For a given harmonic  $h$  of  $f_{alt}$ , the function  $F_h(f)$  is intended to have a large value for a frequency that corresponds to a activity-modulated carrier. We compute this score as

$$F_h(f) = \prod_i F_{i,h}(f) \quad (1)$$

where  $F_{i,h}(f)$  is a sub-score for the  $i$ -th recorded spectrum ( $i$ -th  $f_{alt}$ ). This subscore is computed as

$$F_{i,h}(f) = \frac{SP_i(f + h \cdot f_{alt_i})}{\frac{1}{N-1} \sum_{j \neq i} SP_j(f + h \cdot f_{alt_j})}. \quad (2)$$


This function first appropriately shifts  $SP_i(f)$ , the spectrum captured with the microbenchmark active at alternation frequency  $f_{alt_i}$ , so a side-band signal at  $f_c + h \cdot f_{alt_i}$  gives a peak in  $F_{i,h}$  at the carrier frequency  $f = f_c$ , i.e. we score the side-band signals, but the sub-score is “reported” at  $f_c$ .

The value of the sub-score is computed by normalizing the strength of the side-band signal in this spectrum by the average of the other  $N-1$   $f_{alt_j}$  spectra. For side-band signals that do shift in frequency as  $f_{alt}$  changes, the sub-score for a particular  $i$  will be larger than 1 because the side-band signal is stronger at the  $f_c + f_{alt}$  frequency in this spectrum. At the exact same frequency in at least some of the other spectra, however, the signal will not be as strong because these spectra have peaks at  $f_{alt_j}$  and so their side-band signal is at a different frequency. In contrast, a strong signal that does not shift in frequency as  $f_{alt}$  changes will stay at the same frequency in the other spectra, so the normalization will produce a score close to 1. The overall score  $F_h(f)$  multiplies the sub-scores, so the overall score is close to 1 if no  $f_{alt}$ -induced frequency shifting occurs. If each  $i$ -th spectrum has side-band signals at  $f_{alt_i}$ , the frequency-shifted sub-scores will align producing a very large value for the carrier frequency. Finally, if only some side-band signals are present (one or a few may be “buried” by some unrelated signal), the overall score will be weakened because each obscured  $f_{alt_i}$  side-band will have a sub-score close to 1, but the remaining sub-scores will still increase the overall score significantly above 1. Overall, this heuristic produces large peaks at frequencies of modulated carriers and is almost completely flat at all other frequencies. Figure 9 shows the heuristic function’s output for the carriers shown in Figures 7 and 12.

### 3. Experimental Setup

We evaluate the effectiveness of our FASE methodology by applying it to several computer systems. Unless otherwise indicated, the EM emanations were received with a magnetic loop antenna (AOR LA400) from a distance of 30 cm and





**Figure 9: Output of the heuristic for the 1st and -1st harmonics of  $f_{alt}$  for two carriers.**

a spectrum analyzer (Agilent MXA N9020A) was used to record the spectra of the received signals. This setup was used because it allowed us to capture emanations from the entire system across a wide range of frequencies with little manual effort. We note, however, that attacks exploiting a particular set of carrier signals could likely be carried out at larger distances using more directive antennae optimized for higher gain across a narrower frequency band.

We performed three measurement campaigns, each across a different frequency range and with different FASE parameters, as shown in Figure 10. Parameters  $f_{alt1}$  and  $f_{\Delta}$  were chosen to ensure sufficient separation between side-band and carrier, and between the peaks generated at  $f_{alt1}$ ,  $f_{alt2}$ , etc. Aside from this consideration, the choice of  $f_{alt1}$  and  $f_{\Delta}$  is arbitrary, with the caveat that while using only one choice of  $f_{alt1}$  and  $f_{\Delta}$  is almost always sufficient to detect all carriers, measuring with multiple choices of  $f_{alt1}$  and  $f_{\Delta}$  increases the confidence that all carriers have been detected. For example, a carrier might be missed if FASE is only run with one choice of  $f_{alt1}$  and  $f_{\Delta}$  and a carrier is weak and strong signals happen to occur at the side-bands. We found that five alternation frequencies (i.e.  $f_{alt1}$  through  $f_{alt1} + 4f_{\Delta}$ ) are sufficient to detect almost any carrier even in the presence of unrelated signals from other system activity, noise, and radio broadcasts.

The  $f_{res}$  parameter is the resolution of spectrum sampling. For example, our 0-4MHz measurements used  $f_{res} = 50\text{Hz}$ , so each recorded spectrum has  $4\text{MHz}/50\text{Hz} = 80,000$  data points (frequencies). Each spectrum was measured 4 times over several hours and averaged, and we used the heuristic function in Section 2.4 to detect the 1st, 2nd, 3rd, 4th and 5th positive and negative harmonics of the alternation activity. We then visually inspected the heuristic function’s output to identify peaks (potential carriers). Algorithms to detect peaks in the output of the heuristic function are beyond the scope of this paper ([29] and [4] cover such algorithms), but we

Frequency Range(MHz)	$f_{res}(\text{Hz})$	$f_{alt1}(\text{kHz})$	$f_{\Delta}(\text{kHz})$
0 to 4	50	43.3	0.5
0 to 120	500	43.3	5.0
0 to 1200	500	1800	100

**Figure 10: FASE measurement parameters.**

found that the heuristic function’s output had strong spikes for carriers modulated by system activity, so the task of visually inspecting the output to identify potential carriers was relatively straightforward and quick.


A variety of activities were used as activities X and Y in the alternation loop (Section 2.2) – integer multiplication, division, addition, subtraction, as well as load and store to all levels of the cache hierarchy. The results we show focus on only three X/Y alternations. The first alternates between a load from main memory (LLC miss) and a load from L1 cache (L1 hit), which we abbreviate as LDM/LDL1. This alternation is useful in exposing modulated carriers related to memory activity. We tried other X/Y activity pairs that included main-memory accesses and on-chip activity, e.g. LDM/ADD, LDM/DIV, etc. and also pairings that used STM (LLC write-back activity) instead of LDM. We found that they have some variations in the exact shape and strength of the side-band signals, but applying FASE to them exposes the same carriers as LDM/LDL1.

The second X/Y alternation whose results we show alternates between L2 hits and L1 hits (LDL2/LDL1). This alternation is useful in exposing carriers related to variations in activity on the processor chip. We tried numerous other pairings of on-chip activities, e.g. LDL1/ADD, LDL2/DIV, etc. and found that they expose the same carriers through FASE, although they vary in the exact shape and strength of the side-band signals.

Use of LDM, LDL2 and LDL1 is also methodologically convenient in that it uses the exact same micro-benchmark code for all three activities. They differ only in the `mask` values in Figure 6, which gives us excellent confidence that any observed modulation is due to differences between LDM, LDL1, and LDL2 activity and not the other activity (address computation, looping, etc.) in the alternation loop. Finally, note that the microbenchmarks produce a nearly 100% load, so frequency scaling does not affect our experiments much. However, the effect of frequency scaling wasn’t of interest for these measurements and so we disabled dynamic frequency scaling whenever possible.

## 4. Experimental Results

We begin with Figure 11, which shows the FASE results for a recent desktop system with an Intel Core i7 processor, with the memory access modulation (LDM/LDL1) micro-benchmark. To emphasize the usefulness of FASE, we show a light gray outline of the actual recorded spectrum for one of the alternation frequencies. This spectrum is very noisy and crowded, especially in the long-wave (30-300kHz) and AM radio (540-1600kHz) bands, but FASE correctly indicates which signals



**Figure 11: FASE results for the Intel Core i7 desktop and main-memory (LDM/LDL1) modulating activity.**

are AM-modulated by the alternation activity. The thick vertical lines correspond to the frequency and magnitude of the modulated carrier signals automatically identified by FASE. Lines with the same color/pattern correspond to harmonics of the same frequency. A set of harmonics is likely caused by a periodic yet non-sinusoidal behavior within the system, and the magnitudes of the harmonics in a set give us important clues for identifying the source of that carrier signal. Therefore, after performing FASE it is useful to group the identified carriers into sets such that all the carriers within a set occur at frequencies which appear to be multiplies of one another.

The remainder of this section discusses how we used the information provided by FASE such as carrier frequency, harmonics, modulation depth, and modulation activity (e.g. on-chip activity or memory activity) to identify the sources of three types of carrier signals. In the systems not shown, similar types of signals were detected.

#### 4.1. Switching Voltage Regulators

The set of carriers indicated by red dashed lines in Figure 11 occurs at frequencies 315kHz, 630kHz, 945kHz, etc., which are all multiples of 315kHz. Because the even harmonics of this carrier are relatively strong we can conclude that these carriers are likely caused by some behavior that repeats at 315kHz and has a small duty cycle. It is also helpful to look at each harmonic’s shape in the spectrum. While this figure does not provide enough detail to see each harmonic’s shape distinctly, the shape is very similar to that shown in Figure 12 (this figure corresponds to a different regulator in the same system). The carrier’s energy is spread around its central frequency by what looks like a Gaussian distribution. Clock signals for digital logic and I/O interfaces (such as memory) are tightly controlled but clocks generated by RC oscillators create carriers like the one in Figure 12.<sup>2</sup>

Switching regulators often use RC oscillators. In computer systems, switching regulators convert the 12V to 24V PSU or battery voltage to 1V to 2V supplies used by processors and memory. The duty cycle of the regulator’s switching signal is

small when the ratio between the input and output voltage is large, which is consistent with the 315kHz signal being related to a switching voltage regulator. We manually localized the source of the signal using an EM probe to determine where the 315kHz EM signal was strongest in the system. We found that the signal was strongest near the high power MOSFET switches and power inductors that supply power to the main memory DIMMs. These switches were driven by a nearby switching voltage regulator IC and its switching frequency was 315kHz, confirming our initial hypothesis.


Once the source was found, the modulation mechanism was obvious: the regulator maintains the voltage supplied to the CPU by varying the duty cycle of the control signal of a switch between the 12V supply and the 1V output supply. For example, when DIMMs draw more current, the voltage at the regulator’s output drops, so the regulator compensates by increasing the duty cycle of the switch, i.e. by connecting the 12V supply to its output for a longer fraction of the fixed 315kHz period. When running the LDM/LDL1 microbenchmark, the DRAM regulator’s duty cycle is increased during the DRAM accesses (LDM) and decreased during L1 cache hit activity (LDL1). Changing the duty cycle changes (modulates) the amplitude of all the signal’s harmonics, so LDM/LDL1 activity modulates the emanated signal at the harmonics of the regulator’s switching frequency.

Carrier signals indicated by black dash-dot lines in Figure 11 are also caused by another voltage regulator. This regulator powers the on-chip memory interface (the chip has separate power supplies for its cores and its memory interface). Figure 13 shows the spectrum for heavy on-chip alternation activity (LDL2/LDL1). Only one type of carrier was found to be modulated in this case – the signal that corresponds to the switching regulator for the CPU cores. Figure 12 shows one of the harmonics of this signal in greater detail. We confirmed the origin of both memory interface and core regulator signals through the same near-field localization process. Interestingly, the prominent Gaussian-like shapes of the core regulator’s signal are also visible in Figure 11 but were not reported by FASE because they were not significantly modulated by the LDM/LDL1 alternation.


In many recent processors, the core CPU voltage is adjusted dynamically, while many on-chip cache and memory interface designs require fixed voltage supplies. Therefore, some processors require separate voltage regulators for the CPU and cache. As we have demonstrated, a regulator’s carrier is modulated by the activity in the circuit it powers, so an attacker can distinguish cache and CPU activity by demodulating each regulator’s carrier separately. Also, when separate dynamic voltage scaling is used for each CPU core, each core requires a separate regulator. When such regulator switching frequencies are not identical, attackers might be able to remotely receive a separate power consumption readout for each core, allowing attackers to remotely perform a separate power analysis attack for each core.

<sup>2</sup>This variation (called jitter or phase noise) is well studied because it impacts reliable communications and high frequency digital circuits [21, 35].





**Figure 12: A switching regulator related carrier at  $f_c$  and its right and left side-bands generated by on-chip activity.**



**Figure 13: FASE results for Intel Core i7 desktop and L2 cache (LDL2/LDL1) modulating activity.**

Finally, we note that the emerging use of in-package/on-chip regulators for processors affects regulator-related EM information leakage in new and interesting ways. On-chip linear regulators [37] do not produce modulated emanations because they have no switching frequencies to modulate. The integration of switching regulators has a more complex impact. Each integrated regulator supplies a smaller part of the chip, so the switching currents are lower and follow shorter paths, reducing emanations. However, integrated switching regulators use higher switching frequencies (e.g. 140MHz in [10]) resulting in stronger emanations. Higher switching frequencies also allow faster reactions to changes in the output voltage providing attackers with a higher bandwidth readout of power consumption.


#### 4.2. Memory Refresh

The modulated carrier shown in Figure 11 as solid blue lines has harmonics at frequencies of 512kHz, 1024kHz, etc. This signal did not match any previously known mechanisms that can cause EM emanations. It has a very stable frequency, indicating it was likely generated by logic that is clocked with a crystal-oscillator derived clock. Its harmonics are all of similar strength, indicating an extremely small (<5%) duty cycle. Localization showed that this signal was strongest near

the memory DIMMs. Additional experiments showed that the carrier signal is strongest when there is no memory activity and weakest when we generate continuous memory activity.

This is unusual – if this signal is caused by memory activity, we would expect it to get stronger with more activity. Further measurements with small probes close to the memory revealed many additional harmonics with a greatest common divisor of 128kHz, not 512kHz. This was the key clue in solving the puzzle, because 128kHz corresponds to a period of 7.8  $\mu$ sec, the maximum allowable average time between refresh commands for recent DRAM standards such as DDR3.

While it would be difficult to conclusively prove that this signal is generated by memory refresh activity, the evidence strongly suggests it is. The duty cycle of the memory refresh activity is very low (< 3%) because each refresh command only lasts approximately 200 nsec and occurs every 7.8  $\mu$ sec. The refresh timing is derived from the memory controller clock, which is crystal-derived. While DRAM standards specify that the average time between refresh commands must not exceed 7.8  $\mu$ sec, the memory controller has some control over the timing of the refresh commands. For example, the memory controller could postpone sending refresh commands during a 40  $\mu$ sec period of intense memory activity, and then “catch up” when memory has some idle time. This explains the strangest observation about this signal, which was that it weakens (instead of getting stronger) as memory activity increases. When the memory is inactive, the memory controller simply sends memory refresh commands at regular intervals, resulting in the strongest signal at that interval’s frequency. As memory activity increases, the memory accesses increasingly interfere with the timing of the refresh commands, causing refreshes to be delayed and disrupting their periodicity (thus spreading their emanated energy across a much larger frequency range and causing the signals at 128kHz, 256kHz, etc. to weaken). Although the first harmonic of this signal is weaker than regulator-related signals, note that memory refresh produces many modulated harmonics and that attackers can potentially correlate them to dramatically improve their




**Figure 14: DRAM clock spectrum with 0% (LDL1/LDL1) and 100% (LDM/LDM) memory activity.**

detection of this signal and its signal-to-noise ratio. It is also worth noting that since refresh timing is dictated by a standard, refresh carrier signals are present at roughly the same frequencies on all the systems we tested, which could simplify the exploitation of this leakage.


Detailed exploration of alleviating solutions is beyond the scope of this paper, but this potential problem likely has an easy fix: randomizing the issue of memory refresh commands would be compatible with existing DRAM standards and would greatly reduce the modulation of refresh activity.

#### 4.3. DRAM Memory Clock

Above 30MHz, electromagnetic compatibility (EMC) standards limit the allowable level of EM emanations from consumer devices such as computers. Many periodic signals such as high frequency processor and memory clocks are strong enough to violate these limits, so alleviation techniques for these clock signals have been developed. EMC requirements specify the maximum magnitude for emissions at any particular frequency, and a popular technique (called spread spectrum clocking) varies the clock frequency periodically, spreading the emitted energy across a range of frequencies (instead of emanating it all at one frequency). For example, a 333MHz memory clock might be swept back and forth between 332MHz and 333MHz over a period of 100  $\mu$ sec, producing a spectrum similar to Figure 14. While such techniques facilitate compliance, the signals are only weaker in an averaged sense: attackers can still track the carrier and use the full power of the signal after demodulation. Such “carrier tracking” techniques have already been developed in telecommunications to allow reception of



**Figure 15: DRAM clock spectrum with 50% (LDM/LDL1) mem-ory activity.**



**Figure 16: Heuristic carrier detection function output.**


radio signals transmitted using this technique [13]. Therefore, predictable spread-spectrum clocking does not mitigate information leakage, but it does create interesting problems for discovering such modulated carriers through manual analysis of the spectrum. The shape of the carrier and its side-bands is less recognizable, and the carrier and its side-band signals are likely to overlap significantly when using modulation activity that is not carefully chosen.

To allow FASE to successfully detect modulated spread-spectrum clocks, it is best to set  $f_{alt}$  large enough to move the side-band signals outside of the carrier’s own spectrum. Figure 15 shows the effect of modulating the clock signal at several such alternation frequencies, and Figure 16 illustrates that the heuristic function does detect such modulated signals though it reports the clock as two separate carriers at the edges of the spread out clock signal.

#### 4.4. Testing the Laptop Systems

We tested three laptop systems: one based on an Intel Core i3 processor from 2010, one based on AMD Turion X2 from 2007, and one based on Intel Pentium 3M from 2002. In all three systems, FASE finds the same types of carriers we already reported: regulator-related signals, signals caused by memory refresh, and DRAM clock signals. For example, Figure 17 shows the modulated carrier signals found for the AMD Turion X2 system with LDM/LDL1 alternation of activity. Interestingly, the memory refresh carrier for the AMD Turion X2 laptop is at 132kHz instead of 128kHz as observed in all three other systems. We also confirmed a memory regulator carrier and while the two signals shown as “unidentified” appear to be caused by regulators, we did not confirm their sources because the laptop is very compact and taking it apart to perform localization may damage the system.

The AMD system was the only system confirmed to have an activity-modulated carrier that is not reported by FASE. This carrier was emanated by the voltage regulator circuitry for the processor cores, and was *frequency*-modulated (we confirmed this with a spectrogram of the modulation). Therefore FASE correctly does not report it. This particular regulator keeps the input-to-output switch turned on for a fixed amount of time during its switching cycle, but changes the duration of the switching cycle (i.e. its switching frequency) to increase/decrease its duty cycle. In principle, signals that are



**Figure 17: FASE results for the AMD Turion X2 laptop and main-memory (LDM/LDL1) modulating activity.**

frequency-modulated by system activity should be possible to identify by a FASE-like approach based on spectral properties of FM-modulated signals.

## 5. Related Work

Many side channels have been reported, such as those that observe power consumption [8, 20, 25, 27], sound [5, 12, 33], electromagnetic (EM) emanations [2, 17, 24], behavior under faults [9, 19], and performance of shared caches [6, 36, 38], instruction caches, branch predictors [1], etc.

EM emanations that potentially leak sensitive information are often more easily identified and analyzed in the frequency domain (spectrum) than in the time domain [16, 34]. Some reported EM-emanations attacks rely on AM demodulation of signals [3, 28, 31], but explanations of how the carrier frequency used in the attack was found are either omitted or brief, and no systematic procedures for finding carrier signals modulated by specific system activities have been reported.

Other works focus on finding frequencies where cryptographic keys are leaked by cryptographic algorithms [7], and information theory has been used to correlate observed signals to secret keys [18]. However, these approaches do not provide insight into which aspect of processor, memory, or other system activity modulates the signal at that frequency. Therefore, these methods can be used to evaluate the security of systems against specific attacks on specific cryptographic implementations, but they cannot be used to directly mitigate information leakage. Algorithms have been developed for detecting modulated signals [14] often in the context of military communications intelligence. While such algorithms may discover the same signals FASE does, they would also report radio stations and other modulated signals that are unrelated to the system activity of interest.

Modulated EM signals have been deliberately generated by program activity to measure their propagation distance [39] (distances of at least 2-3m have been reported). The closest work to ours is probably the use of similar microbenchmarks to generate simple (unmodulated) periodic signals at a desired frequency in order to enable measurement of overall poten-

tial for information leakage at the instruction level [11]. In contrast, we use micro-benchmark activity to modulate existing periodic signals to enable automatic identification of such activity-modulated periodic (carrier) signals.

Finally, many of the emanations that FASE reports as activity-modulated signals, such as those from switching regulators and memory clocks, are also considered in the context of electromagnetic compatibility (EMC) and interference (EMI) compliance testing [23, 30]. Work directed at understanding EMI sources often yields useful information about potential carriers of sensitive information. For example, it has been shown that the magnitude of emissions from DRAM modules changes with program activity [26], implicitly indicating that the memory clock acts as a carrier and can be modulated.

## 6. Conclusion

Efficient “surgical” mitigation of side-channel vulnerabilities requires finding information-leaking signals and determining how information is embedded into these signals. In this paper we describe FASE, a novel methodology for automatically finding which EM-emanated signals from a computer system are amplitude-modulated by specific program activities. FASE uses microbenchmarks to generate detectable spectral patterns in the side-bands of all the carrier signals that are AM-modulated by specific system activities, automatically processes measured spectra to identify these patterns, and calculates the frequencies of the modulated carriers.

This approach has several advantages. First, it directly identifies the carrier frequencies modulated by *specific system activities*, which goes a long way toward determining the sources of compromising emanations. Second, it is robust against the interference of unmodulated signals and noise inside and outside of the system, such as AM-modulated signals and carrier-like signals which are not specifically modulated by system activity. Third, it quantifies how strongly carrier signals are modulated, which is useful for identifying how the carrier is generated, for quantifying information leakage, and for evaluating the effectiveness of mitigation efforts. Fourth, it is specifically designed to robustly detect unintentionally modulated signals, which have several inconvenient features not found in ideal AM signals. Finally, each FASE evaluation requires only a few spectrum measurements while other techniques such as DPA require thousands of spectrum captures with different keys and plaintexts [34].

To demonstrate FASE’s effectiveness, we applied it to several computer systems and found activity-modulated signals generated by voltage regulators, memory refresh activity, and DRAM clocks. Our results indicate that separate signals may carry different information about system activity, potentially enhancing an attacker’s capability to extract sensitive information. We also confirm that our methodology correctly separates emanated signals that are affected by specific processor and/or memory activity from those that are not.

## References

- [1] O. Aciğmez, c. K. Koç, and J.-P. Seifert, "On the power of simple branch prediction analysis," in *Proceedings of the 2nd ACM Symposium on Information, Computer and Communications Security*. ACM Press, Mar. 2007, pp. 312–320.
- [2] D. Agrawal, B. Archambeault, J. R. Rao, and P. Rohatgi, "The EM side-channel(s)," in *Proceedings of the Workshop on Cryptographic Hardware and Embedded Systems (CHES)*, 2002, pp. 29–45.
- [3] —, "The EM side-channel(s): attacks and assessment methodologies," 2002. [Online]. Available: <http://www.research.ibm.com/intsec/emf-paper.ps>
- [4] Z. Alfassi, *Statistical Treatment of Analytical Data*. Wiley, 2009.
- [5] M. Backes, M. Durmuth, S. Gerling, M. Pinkal, and C. Spöck, "Acoustic side-channel attacks on printers," in *Proceedings of the USENIX Security Symposium*, 2010.
- [6] E. Bangerter, D. Gullasch, and S. Krenn, "Cache games - bringing access-based cache attacks on AES to practice," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2011.
- [7] A. Barengi, G. Pelosi, and Y. Tegli, "Information Leakage Discovery Techniques to Enhance Secure Chip Design," in *Information Security Theory and Practice. Security and Privacy of Mobile Devices in Wireless Communication*, C. Ardagna and J. Zhou, Eds. Springer Berlin Heidelberg, 2011, vol. 6633, pp. 128–143.
- [8] A. G. Bayrak, F. Regazzoni, P. Brisk, F.-X. Standaert, and P. Ienne, "A first step towards automatic application of power analysis countermeasures," in *Proceedings of the 48th Design Automation Conference*, 2011.
- [9] E. Biham and A. Shamir, "Differential cryptanalysis of the data encryption standard," in *Proceedings of the International Cryptology Conference*, 1993.
- [10] E. A. Burton, G. Schrom, F. Paillet, J. Douglas, W. J. Lambert, K. Radhakrishnan, and M. J. Hill, "FIVR—Fully integrated voltage regulators on 4th generation Intel Core SoCs," in *Applied Power Electronics Conference and Exposition (APEC), 2014 Twenty-Ninth Annual IEEE*. IEEE, 2014, pp. 432–439.
- [11] R. Callan, A. Zajic, and M. Prvulovic, "A practical methodology for measuring the side-channel signal available to the attacker for instruction-level events," in *Proceedings of the 47th International Symposium on Microarchitecture*, 2014.
- [12] S. Chari, J. R. Rao, and P. Rohatgi, "Template attacks," in *Proceedings of the Workshop on Cryptographic Hardware and Embedded Systems (CHES)*, 2002, pp. 13–28.
- [13] B.-Y. Chung, C. Chien, H. Samuelli, and R. Jain, "Performance analysis of an all-digital BPSK direct-sequence spread-spectrum IF receiver architecture," *Selected Areas in Communications, IEEE Journal on*, vol. 11, no. 7, pp. 1096–1107, Sep 1993.
- [14] O. A. Dobre, A. Abdi, Y. Bar-Ness, and W. Su, "Survey of automatic modulation classification techniques: classical approaches and new trends," *Communications, IET*, vol. 1, no. 2, pp. 137–156, 2007.
- [15] R. Erickson and D. Maksimovic, *Fundamentals of Power Electronics*. Springer, 2001.
- [16] C. Gebotys, S. Ho, and C. Tiu, "EM Analysis of Rijndael and ECC on a Wireless Java-Based PDA," in *Cryptographic Hardware and Embedded Systems – CHES 2005*, J. Rao and B. Sunar, Eds. Springer Berlin Heidelberg, 2005, vol. 3659, pp. 250–264.
- [17] D. Genkin, I. Pipman, and E. Tromer, "Get Your Hands Off My Laptop: Physical Side-Channel Key-Extraction Attacks on PCs," in *Proceedings of the Workshop on Cryptographic Hardware and Embedded Systems (CHES)*, 2014.
- [18] B. Gierlichs, L. Batina, P. Tuyls, and B. Preneel, "Mutual information analysis," in *Cryptographic Hardware and Embedded Systems—CHES 2008*. Springer, 2008, pp. 426–442.
- [19] C. Giraud, "DFA on AES," in *Proceedings of the 4th International AES Conference*. Springer, 2003, pp. 27–41.
- [20] L. Goubin and J. Patarin, "DES and Differential power analysis (the duplication method)," in *Proceedings of the Workshop on Cryptographic Hardware and Embedded Systems (CHES)*, 1999, pp. 158–172.
- [21] A. Hajimiri and T. Lee, "A general theory of phase noise in electrical oscillators," *Solid-State Circuits, IEEE Journal of*, vol. 33, no. 2, pp. 179–194, Feb 1998.
- [22] K. B. Hardin, J. T. Fessler, and D. R. Bush, "Spread spectrum clock generation for the reduction of radiated emissions," in *Electromagnetic Compatibility, 1994. Symposium Record. Compatibility in the Loop., IEEE International Symposium on*. IEEE, 1994, pp. 227–231.
- [23] Henry W. Ott, *Electromagnetic Compatibility Engineering*. Wiley, 2009.
- [24] M. G. Khun, "Compromising emanations: eavesdropping risks of computer displays," *The complete unofficial TEMPEST web page*: <http://www.eskimo.com/~joeln/tempest.html>, 2003.
- [25] P. Kocher, J. Jaffe, and B. Jun, "Differential power analysis: leaking secrets," in *Proceedings of the International Cryptology Conference*, 1999, pp. 388–397.
- [26] P. Lee, J. Lee, D.-k. Yoon, J. Choi, and S. Hong, "Analysis of DRAM EMI dependence on data pattern and power delivery design using a near-field EMI scanner," in *Electromagnetic Compatibility and 19th International Zurich Symposium on Electromagnetic Compatibility, 2008. APEMC 2008. Asia-Pacific Symposium on*. IEEE, 2008, pp. 271–274.
- [27] T. S. Messergers, E. A. Dabbish, and R. H. Sloan, "Power analysis attacks of modular exponentiation in smart cards," in *Proceedings of the Workshop on Cryptographic Hardware and Embedded Systems (CHES)*, 1999, pp. 144–157.
- [28] O. Meynard, D. Réal, F. Flament, S. Guilley, N. Homma, and J.-L. Danger, "Enhancement of simple electro-magnetic attacks by pre-characterization in frequency domain and demodulation techniques," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2011. IEEE, 2011, pp. 1–6.
- [29] G. Palshikar, "Simple algorithms for peak detection in time-series," in *Proc. 1st International Conference Advanced Data Analysis, Business Analytics and Intelligence*, 2009.
- [30] C. R. Paul, *Introduction to Electromagnetic Compatibility*, 2nd ed. Wiley, 2006.
- [31] G. Perin, L. Torres, P. Benoit, and P. Maurine, "Amplitude demodulation-based EM analysis of different RSA implementations," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2012. IEEE, 2012, pp. 1167–1172.
- [32] T. Rappaport, *Wireless Communications: Principles and Practice*. Dorling Kindersley, 2009.
- [33] A. Shamir and E. Tromer, "Acoustic cryptanalysis (On nosy people and noisy machines)," <http://tau.ac.il/~tromer/acoustic/>.
- [34] T. Sugawara, Y.-i. Hayashi, N. Homma, T. Mizuki, T. Aoki, H. Sone, and A. Satoh, "Spectrum analysis on cryptographic modules to counteract side-channel attacks," in *EMC*, vol. 9, 2009, pp. 21–24.
- [35] P. Trischitta and E. Varma, *Jitter in Digital Transmission Systems*. Artech House, 1989.
- [36] Y. Tsunoo, E. Tsujihara, K. Minematsu, and H. Miyauchi, "Cryptanalysis of block ciphers implemented on computers with cache," in *Proceedings of the International Symposium on Information Theory and its Applications*, 2002, pp. 803–806.
- [37] Y. Wang and D. Ma, "Ultra-fast on-chip load-current adaptive linear regulator for switch mode power supply load transient enhancement," in *Applied Power Electronics Conference and Exposition (APEC), 2013 Twenty-Eighth Annual IEEE*, March 2013, pp. 1366–1369.
- [38] Z. Wang and R. B. Lee, "New cache designs for thwarting software cache-based side channel attacks," in *Proceedings of the 34th International Symposium on Computer Architecture (ISCA)*. ACM, 2007, pp. 494–505.
- [39] A. Zajic and M. Prvulovic, "Experimental demonstration of electromagnetic information leakage from modern processor-memory systems," *IEEE Transactions on Electromagnetic Compatibility*, vol. 99, no. 3, pp. 1–9, 2014.